AFRL-OSR-VA-TR-2014-0236

Theoretic Methods in Simulation Optimization

Peter Frazier
CORNELL UNIVERSITY

09/24/2014
Final Report

Air Force Research Laboratory
AF Office Of Scientific Research (AFOSR)/ RTA
Arlington, Virginia 22203
Air Force Materiel Command

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

1. REPORT DATE *(DD-MM-YYYY)*	**2. REPORT TYPE**	**3. DATES COVERED** *(From - To)*
14-09-2014	Final	15-06-2011 - 14-06-2014

4. TITLE AND SUBTITLE
Decision-theoretic Methods in Simulation Optimization

5a. CONTRACT NUMBER
FA9550-11-1-0083

5b. GRANT NUMBER
AFOSR

5c. PROGRAM ELEMENT NUMBER

6. AUTHOR(S)
Dr. Peter I. Frazier

5d. PROJECT NUMBER

5e. TASK NUMBER

5f. WORK UNIT NUMBER

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)
Cornell University
School of Operations Research and Information Engineering
Ithaca, NY 14853

8. PERFORMING ORGANIZATION REPORT NUMBER

9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES)
Air Force Office of Scientific Research
875 North Randolph Street
Arlington VA 22203

10. SPONSOR/MONITOR'S ACRONYM(S)
AFOSR

11. SPONSOR/MONITOR'S REPORT NUMBER(S)

12. DISTRIBUTION / AVAILABILITY STATEMENT

Distribution A -- approved for public release

13. SUPPLEMENTARY NOTES

14. ABSTRACT

Simulation optimization is the process of selecting the best among a collection of options, where each option can only be evaluated through stochastic simulation. A good simulation optimization algorithm enables finding the best option, or a very good option, in a reasonable amount of time. Because time is limited, the quality of the simulation optimization algorithm used can often determine whether we are able to discover a high quality option, or merely a mediocre one. The Department of Defense uses stochastic simulators to make a variety of critical strategic decisions, and good simulation optimization algorithms are critical to this decision-making process. The work in this project uses decision-theory to study and improve the decisions made by simulation optimization algorithms, providing algorithms for a number of distinct simulation optimization problems that improve performance over previous state-of-the-art algorithms.

15. SUBJECT TERMS

16. SECURITY CLASSIFICATION OF:			**17. LIMITATION OF ABSTRACT**	**18. NUMBER OF PAGES**	**19a. NAME OF RESPONSIBLE PERSON**
a. REPORT	**b. ABSTRACT**	**c. THIS PAGE**			**19b. TELEPHONE NUMBER** *(include area code)*

AFOSR Final Performance Report

Project Title:
Decision-Theoretic Methods in Simulation Optimization

Award Number:
FA9550-11-1-0083

Award Period:
June 15, 2011 to June 14, 2014

Program:
Optimization and Discrete Mathematics
Program Manager (2011): Dr. Donald Hearn
Program Manager (2012–2014): Dr. Fariba Fahroo, Opt.DMath@afosr.af.mil

Principal Investigator:
Dr. Peter I. Frazier
School of Operations Research and Information Engineering
232 Rhodes Hall
Cornell University
Ithaca, NY 14853
pf98@cornell.edu
`http://people.orie.cornell.edu/pfrazier/`

Summary

Simulation optimization is the process of selecting the best among a collection of options, where each option can only be evaluated through stochastic simulation. A good simulation optimization algorithm enables finding the best option, or a very good option, in a reasonable amount of time. Because time is limited, the quality of the simulation optimization algorithm used can often determine whether we are able to discover a high quality option, or merely a mediocre one. The Department of Defense uses stochastic simulators to make a variety of critical strategic decisions, and good simulation optimization algorithms are critical to this decision-making process. The work in this project uses decision-theory to study and improve the decisions made by simulation optimization algorithms, providing algorithms for a number of distinct simulation optimization problems that improve performance over previous state-of-the-art algorithms.

Project Contributions

This project has made contributions to the following distinct simulation optimization problems: ranking and selection; multiple comparisons with a known standard; stochastic root-finding and closely related first-order methods for stochastic convex optimization; and methods for large-scale and continuous noisy global optimization. The project has additionally made contributions to the axiomatic foundations of decision theory for simulation optimization, to general statistical methodology, and to methods for online learning and sequential decision-making under uncertainty, of which simulation optimization algorithms are a special case.

Ranking and Selection

Ranking and selection is perhaps the most fundamental problem in simulation optimization. In this problem, we have a collection of k alternative system configurations, each of which can be simulated. We wish to use simulation, in an efficient manner, to determine which of the systems has the largest underlying expected performance. The focus in the design of ranking and selection algorithms is on sampling efficiently, so as to discover the best more quickly than one could through a naive equal allocation of simulation sampling effort across the alternatives, and also on algorithms that select the best with some pre-specified lower bound on solution quality.

During the project, we have made the following contributions to ranking and selection:

- [Frazier, 2014] develops the first fully sequential elimination procedure for indifference-zone ranking and selection with *tight* worst-case performance bounds. While many other procedures exist for this problem, they all use performance bounds that become extremely loose as the number of alternative systems grow, which leads them to require a very large number of samples to guarantee a pre-specified solution quality in such problems. Our tight lower bound allows this new procedure, called the Bayes-inspired indifference zone (BIZ) procedure, to solve large problems with many fewer

samples. This provides a ranking and selection procedure that both provides a solution guarantee, and can be used for problems with large numbers of alternatives.

Finding a tight lower bound also resolves a long-standing theoretical issue in the simulation literature. This analysis is made possible by a novel link between Bayesian and frequentist analyses. This paper was a finalist in the 2013 INFORMS Junior Faculty Interest Group (JFIG) Paper Competition.

- [Chick and Frazier, 2012] studies the Bayesian ranking and selection problem, in which the goal is to find a procedure with good average-case performance, in contrast with one with a worst-case performance bound. It first studies a simple ranking and selection problem with a single unknown alternative and another known alternative, using a continuous-time approach, and then uses the optimal procedure from this simple problem to devise a procedure for the full ranking and selection problem. In numerical experiments, this new procedure, called the Economics of Sampling (ESP) procedure, outperforms the best previously existing procedures.
- [Xie and Frazier, 2013] derives computable upper bounds on the best performance possible in the Bayesian ranking and selection problem. Computing a Bayes-optimal algorithm is computationally intractable in general, due to the so-called curse of dimensionality, and so having an upper bound on performance is useful because it lets us compute an optimality gap for existing heuristic policies, and gives information about how far from optimal the best existing heuristic is. In problem settings where this optimality gap is small, we can rest assured that the best existing heuristic is doing almost as well as the Bayes-optimal algorithm.
- [Frazier and Kazachkov, 2011] studies a multi-objective formulation of the Bayesian ranking and selection problem, using a novel view of multiobjective optimization, in which the decision-maker has a single objective that combines alternatives' attributes, but this single objective is unknown to the simulation analyst, and will only be revealed after simulation is complete and the decision-maker is shown the result. This matches a common situation in practice, in which eliciting a decision-maker's preferences using hypothetical system performance is more expensive than simply running the simulation and showing estimated performances of real systems. Placing a prior on the decision-maker's preferences then provides a natural algorithm for allocating simulation effort, to explore the most important parts of the Pareto frontier.

Large-scale Bayesian Global Optimization

When optimizing over a set of alternatives parameterized by some multi-dimensional input, e.g., in a continuous space, or in some more complex combinatorial space, the set of alternative systems is typically too large to enumerate, preventing the application of ranking and selection methods. Instead, we must use a simulation optimization algorithm that explicitly models the dependence between alternative systems. Bayesian global optimization offers a natural framework for doing this, allowing a statistical model to be used that models the relationship between alternatives, and also providing a natural notion of value of information that can be used to select points to sample.

This project has made the following contributions to large-scale Bayesian global optimization:

- [Clark, 2012] considers Bayesian global optimization of noise-free continuous functions, when parallel computing resources are available, and we can perform multiple function evaluations at the same time. This PhD thesis derives an algorithm for allocating this batch of measurements that is optimal when just one stage of samples remains. Numerical experiments demonstrate that this method provides a significant speedup over allocating just one sample at a time.

  In a significant transition of this project to industry, the algorithm in this PhD thesis has been adopted by the internet company Yelp.com, which devoted a significant amount of their own resources (half-time for two engineers over the course of one year, and four summer interns) to develop a high-quality open source implementation of this algorithm, available at `http://yelp.github.io/MOE/`. Yelp uses this algorithm for optimizing parameters in the algorithms that run their website. The software they produced is also freely available, and is also being considered for use by Netflix.

- [Scott et al., 2011] provides an easily computed approximation to the knowledge-gradient policy with correlated beliefs for Bayesian global optimization over continuous spaces, which is Bayes-optimal when a single measurement remains. This approximation performs quite well in numerical studies, and enjoys its own convergence guarantees.

- [Frazier et al., 2011, Xie et al., 2014] consider how common random numbers can be used within a Bayesian global optimization algorithm over multi-dimensional discrete spaces, using value of information calculations to choose pairs of points to sample in parallel using common random numbers. Common random numbers induce correlation in the simulation noise, which can reduce the variance in estimates of the difference in quality between pairs of alternatives, and improve overall performance. These papers demonstrate in numerical experiments that the use of common random numbers significantly improves performance over a number of other leading simulation optimization algorithms.

- [Xie et al., 2012] considers a novel form of simulation optimization where we seek to optimize $\int f(x,\omega)p(\omega)d\omega$ over x, and we can evaluate $f(x,\omega)$ at both an x and an ω of our choosing. While this is in contrast with most formulations of the simulation optimization problem, where we choose x but not ω, it appears quite frequently in practice, including an application to cardiovascular bypass graft design considered in this paper. This paper develops a value-of-information method for this problem, and show that it outperforms two other benchmark procedures.

- [Mes et al., 2011, Negoescu et al., 2011] consider two novel forms of discrete optimization via simulation, which arise when searching over spaces that are more complex than $\mathbb{R}^d$. [Negoescu et al., 2011] considers a space of small molecules described by a linear statistical model, such as would be considered in drug discovery, materials science, or chemical engineering, while [Mes et al., 2011] considers a general space of alternatives described by a hierarchical aggregation function, which appears in transportation

and logistics applications, and also a variety of other contexts. Both papers develop value-of-information procedures that are seek to choose the point to sample next that provides the most value in terms of our ability to solve the overall optimization problem.

Stochastic Root-Finding

In stochastic root-finding, we have a monotonic function that can only be measured with noise, and we wish to use these noisy evaluations at adaptively chosen locations to determine where the root of the function lies. This problem appears in simulation optimization when we have a convex function whose gradient can be observed with noise, and we wish to find the root of the gradient, which is a maximum of the convex function. One well-known method for solving stochastic root-finding problems is stochastic approximation, which sometimes works well, but is quite sensitive to the choice of tuning sequence. In work performed in this project, we pursue an alternative more robust method for based on a probabilistic form of bisection, which is motivated by a decision-theoretic analysis. This method is developed in the following papers:

- [Jedynak et al., 2012] studies a problem in which we wish to find an object using a limited number of adaptively chosen queries, and our queries take the form of choosing a set, and asking whether or not the object resides in this set. The answers to our queries are provided with noise. Using the expected entropy of the posterior distribution as our measure of performance, we show that the greedy policy is Bayes-optimal over arbitrary time horizons. This is surprising, because the greedy policy only optimizes for one period into the future, and is also extremely useful computationally, because it is quite easy to compute, in contrast with general dynamic programming methods for computing Bayes-optimal policies, which require extreme amounts of computation.

- [Waeber et al., 2011, Waeber et al., 2013] applies the analysis of [Jedynak et al., 2012] to a stylized version of the stochastic root-finding problem, in which obtaining a noisy observation of the function with constant error probability at a point x is viewed as equivalent to querying a noisy oracle whether the root is in the set $(-\infty, x)$. It then analyzes the greedy policy, which is optimal for an entropy-based criterion, in terms of frequentist absolute error, and establishes an exponential rate of convergence for this algorithm.

  [Waeber et al., 2011] won the Best Student Paper (OR/MS focused) at the 2011 Winter Simulation Conference.

- [Waeber, 2013] builds on the stylized analysis of the stochastic-root finding problem in [Waeber et al., 2011, Waeber et al., 2013] to provide a practical algorithm that resamples at each queried point, to provide a response whose error probability can be bounded. In numerical experiments, this algorithms exhibits greater robustness than stochastic approximation, and a competitive convergence rate.

- [Pallone et al., 2014] considers noise-free root-finding with parallel computing resources, provides a computable Bayes-optimal policy for this problem, and analyses a simpler

policy with competitive performance.

- [Sznitman et al., 2013] applies insights obtained in [Jedynak et al., 2012] to a problem in scanning electron microscopy that is closely related to the stochastic root-finding problem: adaptive edge detection.

- [Han et al., 2014] considers an extension of the model of [Jedynak et al., 2012], in which multiple objects may be present, and the response to a query is the number of objects in the queried set.

Multiple Comparisons with a Known Standard, and Feasibility Determination

In the problem of multiple comparisons with a known standard, we wish to use simulation to determine which alternative systems have mean performance exceeding a known threshold. This is often required when deciding which systems exhibit acceptable performance, or which have a performance that exceeds that of some existing system currently in use. We can view this problem as a form of feasibility determination, where we wish to determine which alternatives meet some performance constraint, and in this form can be a pre-processing step in a simulation optimization procedure.

In fully sequential multiple comparisons with a standard, we allocate our simulation effort adaptively, to enable this allocation to be more efficient by using more simulation samples for those alternatives that are difficult to classify, and less for those that can be classified easily.

This project has made the following contributions to fully sequential multiple comparisons with a known standard:

- [Xie and Frazier, 2013] shows that the Bayes-optimal policy for fully sequential multiple comparisons with a known standard can be computed efficiently, at least when sampling is not limited by a fixed finite horizon, but is instead limited by an economic cost for each sample, and/or a random horizon, as would occur if the desired sampling budget were unknown when sampling started. Although computing Bayes-optimal policies is usually intractable due to the curse of dimensionality, it is made possible in this case by a link to multi-armed bandits.

  This paper won the 2013 INFORMS Computing Society Student Paper Prize, and was a finalist in the 2011 INFORMS Junior Faculty Interest Group (JFIG) Paper Competition.

- [Hu et al., 2014] considers fully sequential multiple comparisons with a known standard in a setting with a known finite horizon, and parallel sampling resources. Using a Lagrangian relaxation, similar to that used in deriving Whittle index policies for the restless multi-armed bandit problem, it provides a computable upper bound on the performance of a Bayes-optimal policy. This upper bound may be used to bound the sub-optimality of existing heuristic procedures. This upper bound also provides a price-directed heuristic policy whose performance is close to the upper bound across a range of problems.

Simulation Screening

Simulations of real-world phenomena often take hundreds or even thousands of input parameters, but many of these input parameters have very little effect on the simulation output. Simulation screening is the process of determining which inputs have a significant effect on simulation output, and is often done as a preprocessing step before performing simulation optimization, removing insignificant inputs to allow optimization to be performed more efficiently.

- [Frazier et al., 2012] considers the well-known sequential bifurcation algorithm for simulation screening, and a particular design choice within this algorithm that arises when deciding how to split groups of inputs when testing them for significance. This paper uses a Bayesian analysis to compute the optimal group splitting decision. Using this optimal group splitting decision is then shown to improve the efficiency of the sequential bifurcation algorithm.

Decision-Theoretic Foundations of Simulation Optimization

The project made the following contributions to the axiomatic foundations of simulation optimization, and general theory for assessing the quality of simulation optimization algorithms:

- [Waeber et al., 2012] studies the axiomatic foundations of simulation optimization, creating convex risk measures for evaluating the performance of algorithms for simulation optimization. This provides a collection of performance measures, which include the Bayesian performance measure, which considers average case performance with respect to a single prior distribution; the worst-case performance over a collection of possible configurations; but also a range of other performance measures that make different tradeoffs between robustness to prior misspecification and average-case performance.
- [Frazier and Powell, 2011] provides general sufficient conditions that can be used to check the consistency of a simulation optimization procedure, i.e., whether it eventually will discover the best option in the limit as the simulation sampling budget grows large. This is a useful condition to check, as lack of consistency is suggestive that a procedure is poorly designed, but checking this condition is often quite difficult. The sufficient condition in this paper can be used to check for consistency more easily, and the paper uses it to show consistency of several existing ranking and selection procedures.
- [Frazier, 2012] is a tutorial paper, given at the Winter Simulation Conference, which provides a clear introduction to the use of decision-theory and Bayesian methods in simulation optimization.

Online Learning and Decision-making under Uncertainty

Simulation optimization is closely linked to online learning and decision-making under uncertainty, in which we are making decisions that provide rewards, while at the same time

learning the underlying probability distributions that govern these rewards. Indeed, designing an algorithm for simulation optimization with optimal or near-optimal average-case performance is a problem in sequential decision-making under uncertainty. The project's focus on simulation optimization made possible the following synergistic contributions to online learning during the award period:

- [Frazier et al., 2014] studies the multi-armed bandit problem and the exploration vs. exploitation tradeoff in a strategic setting, in which a central planner (e.g., a federal funding agency) in faced with a stream of agents (e.g., researchers seeking funding) who are choosing between a variety of options (e.g., research approaches or research problems). The options provide random payoffs coming from unknown probability distributions. Agents share a common prior distribution over the reward distributions, which is updated as the rewards from previous agents are observed. We model these agents as myopic and self-interested, and so left on their own would pursue the option with the largest expected posterior reward, and must be incentivized if they are to more fully explore high-risk high-reward options. This paper characterizes the achievable total expected sum of rewards across all agents, in terms of the budget used to incentivize exploration.

  This paper won the Best Paper Award at the 15th ACM Conference on Economics and Computation, 2014.

- [Ryzhov et al., 2012] provides an easy-to-compute approximation to the Bayes-optimal policy for the multi-armed bandit problem, called the online knowledge-gradient policy. This policy is seen to perform well in comparison with the optimal policy and a variety of other heuristics.

- [Ryzhov et al., 2014] provides a new optimal stepsize rule for use in approximate dynamic programming algorithms, which can be used for solving general sequential decision-making problems.

- [Zhang et al., 2012, Zhao and Frazier, 2014] consider applications of decision-making under uncertainty and adaptive information collection, arising in medical surveillance and information retrieval, respectively.

Statistical Methodology

Decision-theoretic simulation optimization involves a great deal of Bayesian statistics. During the project, in addition to making advances to algorithms specifically for simulation optimization, we also made fundamental advances in nonparametric Bayesian statistics itself, which can be applied to simulation optimization, but also a variety of other applications. During the award period, we made the following synergistic statistical contributions:

- [Blei and Frazier, 2011] and [Gershman et al., 2014] consider nonparametric Bayesian mixture models, providing so-called "distance dependent" versions of the well-known Chinese Restaurant Process and Indian Buffet Process (IBP), which can be used in applications for which the assumption of exchangeability (made by the standard CRP and IBP) is inappropriate.

- [Clark et al., 2013] and [Meltzer et al., 2013] consider statistical applications arising in bioinformatics and medical decision-making.

Journal publications supported by this project

[Blei and Frazier, 2011] Blei, D. M. and Frazier, P. I. (2011). Distance dependent chinese restaurant processes. *Journal of Machine Learning Research*, 12:2461–2488.

[Chick and Frazier, 2012] Chick, S. E. and Frazier, P. I. (2012). Sequential sampling for selection with economics of selection procedures. *Management Science*, 58(3):550–569.

[Clark et al., 2013] Clark, S. C., Egan, R., Frazier, P. I., and Wang, Z. (2013). Ale: a generic assembly likelihood evaluation framework for assessing the accuracy of genome and metagenome assemblies. *Bioinformatics*, 29(4):435443.

[Frazier, 2014] Frazier, P. I. (2014). A fully sequential elimination procedure for indifference-zone ranking and selection with tight bounds on probability of correct selection. *Operations Research*, 62(4):926–942.

[Frazier and Powell, 2011] Frazier, P. I. and Powell, W. B. (2011). Consistency of sequential Bayesian sampling policies. *SIAM Journal on Control and Optimization*, 49:712–731.

[Gershman et al., 2014] Gershman, S. J., Frazier, P. I., and Blei, D. M. (2014). Distance dependent infinite latent feature models. *IEEE Transactions on Pattern Analysis and Machine Intelligence*. to appear, published online May 22, 2014.

[Jedynak et al., 2012] Jedynak, B., Frazier, P. I., and Sznitman, R. (2012). Twenty questions with noise: Bayes optimal policies for entropy loss. *Journal of Applied Probability*, 49(1):114–136.

[Lee et al., 2013] Lee, L. H., Chew, E. P., Frazier, P. I., Jia, Q. S., and Chen, C. H. (2013). Advances in simulation optimization and its applications. *IIE Transactions*, 45(7):683–684.

[Meltzer et al., 2013] Meltzer, A. J., Graham, A., Connolly, P. H., Karwowski, J. K., Bush, H. L., Frazier, P. I., and Schneider, D. B. (2013). Risk factors for early failure after peripheral endovascular intervention: Application of a reliability engineering approach. *Annals of Vascular Surgery*, 27(1):53–61.

[Mes et al., 2011] Mes, M. R. K., Powell, W. B., and Frazier, P. I. (2011). Hierarchical knowledge gradient for sequential sampling. *Journal of Machine Learning Research*, 12(Oct):2931–2974.

[Negoescu et al., 2011] Negoescu, D. M., Frazier, P., and Powell, W. B. (2011). The knowledge gradient algorithm for sequencing experiments in drug discovery. *INFORMS Journal on Computing*, 23(1).

[Ryzhov et al., 2014] Ryzhov, I., Frazier, P. I., and Powell, W. B. (2014). A new optimal stepsize rule for approximate dynamic programming. IEEE Transactions on Automatic Control, to appear.

[Ryzhov et al., 2012] Ryzhov, I. O., Powell, W. B., and Frazier, P. I. (2012). The knowledge gradient algorithm for a general class of online learning problems. *Operations Research*, 60(1).

[Scott et al., 2011] Scott, W., Frazier, P. I., and Powell, W. B. (2011). The correlated knowledge gradient for simulation optimization of continuous parameters using gaussian process regression. *SIAM Journal on Optimization*, 21(3):996–1026.

[Waeber et al., 2012] Waeber, R., Frazier, P. I., and Henderson, S. G. (2012). A framework for selecting a selection procedure. *ACM Transactions on Modeling and Computer Simulation*, 22(3).

[Waeber et al., 2013] Waeber, R., Frazier, P. I., and Henderson, S. G. (2013). Bisection search with noisy responses. *SIAM Journal on Control and Optimization*, 51(3):2261–2279.

[Xie and Frazier, 2013] Xie, J. and Frazier, P. I. (2013). Sequential bayes-optimal policies for multiple comparisons with a known standard. *Operations Research*, 61(5):1174–1189.

Journal papers in review supported by this project

[Han et al., 2014] Han, W., Frazier, P. I., and Jedynak, B. M. (2014). Twenty questions for localizing multiple objects by counting: Bayes optimal policies for entropy loss. in review, arXiv preprint arXiv:1407.4446.

[Xie et al., 2014] Xie, J., Frazier, P. I., and Chick, S. E. (2014). Bayesian optimization via simulation with pairwise sampling and correlated prior beliefs. in review.

[Zhao and Frazier, 2014] Zhao, X. and Frazier, P. I. (2014). Exploration vs. exploitation in the information filtering problem. in review, arXiv preprint arXiv:1407.8186.

Conference publications supported by this project

[Frazier, 2012] Frazier, P. I. (2012). Tutorial: Optimization via simulation with bayesian statistics and dynamic programming. In Laroque, C., Himmelspach, J., Pasupathy, R., Rose, O., and Uhrmacher, A. M., editors, *Proceedings of the 2012 Winter Simulation Conference Proceedings*, pages 79–94, Piscataway, New Jersey. Institute of Electrical and Electronics Engineers, Inc.

[Frazier et al., 2012] Frazier, P. I., Jedynak, B. M., and Chen, L. (2012). Sequential screening: A bayesian dynamic programming analysis. In Laroque, C., Himmelspach, J., Pasupathy, R., Rose, O., and Uhrmacher, A. M., editors, *Proceedings of the 2012 Winter*

Simulation Conference, pages 555–566, Piscataway, New Jersey. Institute of Electrical and Electronics Engineers, Inc.

[Frazier and Kazachkov, 2011] Frazier, P. I. and Kazachkov, A. M. (2011). Guessing preferences: A new approach to multi-attribute ranking and selection. In Jain, S., Creasey, R. R., Himmelspach, J., White, K. P., , and Fu, M., editors, *Proceedings of the 2011 Winter Simulation Conference*, pages 4324–4336, Piscataway, New Jersey. Institute of Electrical and Electronics Engineers, Inc.

[Frazier et al., 2014] Frazier, P. I., Kempe, D., Kleinberg, J., and Kleinberg, R. (2014). Incentivizing exploration. In *Proceedings of the 15th ACM Conference on Economics and Computation (EC'14)*, pages 5–22, New York, NY, USA. Association for Computing Machinery.

[Frazier et al., 2011] Frazier, P. I., Xie, J., and Chick, S. E. (2011). Value of information methods for pairwise sampling with correlations. In Jain, S., Creasey, R. R., Himmelspach, J., White, K. P., , and Fu, M., editors, *Proceedings of the 2011 Winter Simulation Conference*, pages 3979–3991, Piscataway, New Jersey. Institute of Electrical and Electronics Engineers, Inc.

[Hu et al., 2014] Hu, W., Frazier, P. I., and Xie, J. (2014). Parallel bayesian policies for finite-horizon multiple comparisons with a known standard. In Tolk, A., Diallo, S. D., Ryzhov, I. O., Yilmaz, L., Buckley, S., and Miller, J. A., editors, *Proceedings of the 2014 Winter Simulation Conference*, Piscataway, New Jersey. Institute of Electrical and Electronics Engineers, Inc.

[Pallone et al., 2014] Pallone, S., Frazier, P. I., and Henderson, S. G. (2014). Multisection: Parallelized bisection. In Tolk, A., Diallo, S. D., Ryzhov, I. O., Yilmaz, L., Buckley, S., and Miller, J. A., editors, *Proceedings of the 2014 Winter Simulation Conference*, Piscataway, New Jersey. Institute of Electrical and Electronics Engineers, Inc.

[Sznitman et al., 2013] Sznitman, R., Lucchi, A., Frazier, P. I., Jedynak, B. M., and Fua, P. (2013). An optimal policy for target localization with application to electron microscopy. In Dasgupta, S. and McAllester, D., editors, *Proceedings of the 30th International Conference on Machine Learning (ICML)*, pages 1–9. Omnipress.

[Waeber et al., 2011] Waeber, R., Frazier, P. I., and Henderson, S. G. (2011). A bayesian approach to stochastic root finding. In Jain, S., Creasey, R. R., Himmelspach, J., White, K. P., , and Fu, M., editors, *Proceedings of the 2011 Winter Simulation Conference*, pages 4038–4050, Piscataway, New Jersey. Institute of Electrical and Electronics Engineers, Inc.

[Xie and Frazier, 2013] Xie, J. and Frazier, P. I. (2013). Upper bounds on the bayes-optimal procedure for ranking & selection with independent normal priors. In Pasupathy, R., Kim, S.-H., Tolk, A., Hill, R., and Kuhl, M. E., editors, *Proceedings of the 2013 Winter Simulation Conference*, pages 877–887, Piscataway, New Jersey. Institute of Electrical and Electronics Engineers, Inc.

[Xie et al., 2012] Xie, J., Frazier, P. I., Sankaran, S., Marsden, A., and Elmohamed, S. (2012). Gaussian-process-based black-box optimization for cardiovascular bypass graft design under parameter uncertainty. In *The 50th Annual Allerton Conference on Communication, Control and Computing*, pages 406–413, Piscataway, New Jersey. Institute of Electrical and Electronics Engineers, Inc.

[Zhang et al., 2012] Zhang, S., Hanagal, P., Frazier, P. I., Meltzer, A. J., and Schneider, D. B. (2012). Optimal patient-specific post-operative surveillance for vascular surgery. In *Proceedings of the 7th INFORMS Workshop on Data Mining and Health Informatics*.

PhD theses

[Clark, 2012] Clark, S. C. (2012). *Parallel Machine Learning Algorithms in Bioinformatics and Global Optimization*. PhD thesis, Cornell University. https://github.com/sc932/thesis.

[Waeber, 2013] Waeber, R. (2013). *Probabilistic Bisection Search for Stochastic Root-Finding*. PhD thesis, Cornell University. http://hdl.handle.net/1813/33929.

[Xie, 2014] Xie, J. (2014). *Bayesian Designs for Sequential Learning Problems*. PhD thesis, Cornell University. http://hdl.handle.net/1813/37171.

Personnel:

PI: Peter I. Frazier

Cornell PhD students supported:
Stephen Pallone, Cornell PhD student.
Rolf Waeber, Cornell PhD 2013
Jialei Wang, Cornell PhD student.
Jing Xie, Cornell PhD 2014

Cornell PhD students involved but not supported:
Scott Clark, Cornell PhD 2012
Shanshan Zhang, Cornell PhD 2013
Xiaoting Zhao, Cornell PhD student

Cornell Masters in Engineering students involved but not supported:
Pu Yang
Saurabh Arora

Cornell Undergraduates involved but not supported:
Darlin Alberto
Saketh Are
Max Berman
John (Massey) Cashore
Michael Dezube

Rutna Gadh
Pranav Hanagal
Christina Hardin
Lemuel Kumarga
Junbao Li
Guangyu (Paul) Liu
Yujia Liu
Divya Singhvi
Somya Singhvi
Andrew Sung
Arya Viswanathan
Li Wang

Awards/Honors (2011-2014)

- Best Paper Award, 15th ACM Conference on Economics and Computation, 2014.
  for P.I. Frazier, D. Kempe, J. Kleinberg, R. Kleinberg, "Inventizing Exploration."

- Finalist, INFORMS Data Mining Best Student Paper, 2014
  for J. Wang, M.D. Burkart, P.I. Frazier, N.Gianneschi, M.K. Gilson, N. Kosa, L. Tallorin, and P. Yang, "Finding Short Peptide Substrates using Bayesian Active Learning."

- Finalist, INFORMS Junior Faculty Interest Group (JFIG) Paper Competition, 2013.
  for P.I. Frazier, "A Fully Sequential Elimination Procedure for Indifference-Zone Ranking and Selection with Tight Bounds on Probability of Correct Selection."

- INFORMS Computing Society Student Paper Prize, 2013.
  for J. Xie, P.I. Frazier, "Sequential Bayes-Optimal Policies for Multiple Comparisons with a Control."

- NSF CAREER Award, 2012, for "Methodology for Optimization via Simulation: Bayesian Methods, Frequentist Guarantees, and Applications to Cardiovascular Medicine," P. I. Frazier

- Best Student Paper (OR/MS focused), Winter Simulation Conference, 2011.
  for R. Waeber, P.I. Frazier, and S.G. Henderson, "A Bayesian Approach to Stochastic Root-Finding."

- Finalist, INFORMS Junior Faculty Interest Group (JFIG) Paper Competition, 2011.
  for J. Xie, P.I. Frazier, "Sequential Bayes-Optimal Policies for Multiple Comparisons with a Control."

Interactions/Transitions/Transfers

Algorithms Transitioned to Industry

Metrics Optimization Engine (MOE), `https://github.com/Yelp/MOE`

MOE is a Bayesian global optimization engine for real-world metric optimization, where a "metric" is understood to be any performance measure. It is used internally by the tech startup Yelp (`http://yelp.com`) to optimize tunable parameters in algorithms that deliver content to users. Because it is open-source, it can also be used and extended by others for any metric optimization applications.

MOE is based on an algorithm developed in this project and described in [Clark, 2012], and was developed by engineers at Yelp: Scott Clark (Yelp, Cornell PhD 2012), Eric Liu (Yelp), Deniz Oktay (Yelp & MIT), Norases Vesdapunt (Yelp & Stanford), Jialei Wang (Yelp & Cornell).

Yelp has been advertising this engine through blog posts (see http://engineeringblog.yelp.com/2014/07/introducing-moe-metric-optimization-engine-a-new-open-source-machine-learning-service-for-optimal-ex.html), and through talks engineers have given at other tech firms in the San Francisco Bay Area, such as Netflix. It has received a great deal of attention from the community, and is already quite popular on the open-source software repository github.com[1]

Consultative and advisory functions to laboratories

- **Air Force Research Lab:**

  Frazier has played a consultative role through email and videoconferences on a project with Benji Maruyama and Daylond Hooper at AFRL, on autonomous systems for growing carbon nanotubes optimally, which uses algorithms based on Frazier's work to decide which experimental conditions to use. Frazier visited AFRL in August 2014, for the Autonomous Research Systems for Materials Development Workshop, organized by Benji Maruyama.

  Frazier has also had a number of conversations (at AFOSR Program Reviews for the Natural Materials and Extremophiles program) and email interchanges with Rajesh Naik at AFRL, on experimental methods for peptide design.

- **Los Alamos National Lab:**

  Frazier visited LANL, hosted by Frank Alexander, in January 2013, where he discussed the use of simulation optimization methods for optimizing large-scale simulations and experimental platforms.

  He also gave a talk through a videoconference link at a workshop involving Frank Alexander, Turab Lookman, and others from LANL, at the Materials Informatics Workshop at the Sante Fe Institute in April 2013. In February 2014, Frazier attended and gave a talk at a follow-on workshop organized by Turab Lookman, the Information

[1] Github usage metrics (updated 9/14/2014): 311 stars, 16 forks, $\geq 6,000$ unique visits.

Science for Materials Discovery and Design Workshop, which was also in Sante Fe, and involved a number of participants from LANL.

- **U.S. Army Training and Doctrine Command Analysis Center TRAC)**

  Frazier has discussed the use of simulation optimization with the US Army Training and Doctrine Command Analysis Center, and specifically with Peter Nesbitt, MAJ U.S. Army, for use with their logistics and resource allocation simulators System for Periodically Apportioning Demands (SPADES) and LBC (logistics battle command). Following an initial round of discussions that occurred remotely through research collaborators Ned Dimitrov (UT Austin) and Dashi Singham (Naval Postgraduate School), Frazier visited Nesbitt at TRAC in Monterey, California in September 2014.

Invited Departmental Seminars and Tutorials

- "Bayesian Methods for Simulation Optimization" Department of Industrial Engineering, Tsinghua University, Beijing, June 2014.

- "Information Filtering for arXiv.org: Bandits, Exploration vs. Exploitation, and the Cold Start Problem" Department of Industrial and Systems Engineering, Georgia Institute of Technology, Atlanta GA, April 2014.

- "Optimal Learning for Discovering Minimal Peptide Substrates." Mirkin Research Group, Department of Chemistry, Northwestern University, Evanston IL, February 2014.

- Masterclass lecture series; "Machine Learning for Optimization" (day 1); "Machine Learning & Sequential Decision-Making" (day 2); "Coding for Researchers" (day 3). STOR-i, Lancaster University, Lancaster UK, January 2013.

- "Information Filtering for arXiv.org." Yelp.com, San Francisco, December, 2013.

- "Bayesian Methods for Simulation Optimization." Center for Information Science and Systems Engineering (CISE), Boston University, November 2013.

- "Bayesian Methods for Simulation Optimization." 3M Research Center, St. Paul MN, October 2013.

- "Bayesian Methods for Simulation Optimization." IBM T.J. Watson Research Center, Yorktown Heights NY, August 2013.

- "Bayesian Methods for Simulation Optimization." Virginia Commonwealth University, Department of Statistical Sciences and Operations Research, January 2013.

- "Tutorial: Optimization via Simulation with Bayesian Statistics and Dynamic Programming," Advanced Tutorials Track, Winter Simulation Conference, Berlin, December 2012.

- “Bayesian Methods for Simulation Optimization.” University of Virginia, Department of Systems and Information Engineering, November 2012.

- “Optimal Learning: an Overview.” The Ohio State University, Department of Psychology, Myung/Pitt Research Group, October 2012.

- “Bayesian Methods for Simulation Optimization.” Rensselaer Polytechnic Institute, Department of Industrial and Systems Engineering, March 2012.

- “Tutorial: Bayesian Methods for Global and Simulation Optimization.” INFORMS tutorial, INFORMS Annual Meeting, Charlotte, November 2011

- “Optimizing Time-consuming Objective Functions: Case Studies in Drug Development and Simulation Calibration.” Department of Energy, Joint Genome Institute, June 2011.

Conference Presentations

- P.I. Frazier, X. Zhao, “Information Filtering for arXiv.org: Bandits, Exploration vs. Exploitation, and the Cold Start Problem” Mostly OM, Beijing, June 2014.

- P.I. Frazier, J. Wang, P. Yang, “Optimal Learning for Discovering Minimal Peptide Substrates.” Information Science for Materials Discovery and Design Workshop, Sante Fe, NM, February 2014.

- P.I. Frazier, R. Waeber, S.G. Henderson, “Bayes-Optimal Methods for Optimization via Simulation: The Probabilistic Bisection Algorithm.” STOR-i Workshop, Lancaster University, Lancaster UK, January 2014.

- P.I. Frazier “Optimal Learning for Peptide Design.” Air Force Office of Scientific Research, Natural Materials, Systems and Extremophiles Program Review, Eglin Air Force Base, Fort Walton Beach, FL, December 2013.

- P.I. Frazier, “A Fully Sequential Elimination Procedure for Indifference-Zone Ranking and Selection with Tight Bounds on Probability of Correct Selection.” INFORMS Annual Meeting, JFIG Prize Presentation, Minneapolis MN, October 2013.

- P.I. Frazier “Ranking and Selection with Tight Bounds on Probability of Correct Selection.” INFORMS Annual Meeting, Minneapolis MN, October 2013.

- P.I. Frazier “Bayesian Active Learning for Finding Maximally-valued Exemplars.” INFORMS Annual Meeting, Minneapolis MN, October 2013.

- P.I. Frazier “Ranking and Selection with Tight Bounds on Probability of Correct Selection.” Applied Probability Society Conference, San Jose, Costa Rica, July 2013.

- P.I. Frazier, S. Zhang, P. Hanagal, A.J. Meltzer, D.B. Schneider, “Optimal Patient-specific Post-operative Surveillance for Vascular Surgeries.” INFORMS Healthcare Conference, Chicago, IL, June 2013.

- P.I. Frazier "Optimal Learning: Bayesian Methods for Simulation Optimization." AFOSR Program Review for Optimization & Discrete Mathematics, Arlington, VA, April 2013.
- P.I. Frazier "Optimal Learning in Materials Science." Materials Informatics Workshop, Sante Fe, NM, April 2013.
- P.I. Frazier "Ranking and Selection with Tight Bounds on Probability of Correct Selection." Simulation Optimization Workshop, Viña del Mar, Chile, March 2013.
- P.I. Frazier, W.B. Powell and H.P. Simão "Simulation Calibration with Correlated Knowledge Gradients." SIAM Conference on Computational Science & Engineering, Bostan, MA, February 2013.
- W.B. Powell and P.I. Frazier "Optimal Learning for Efficient Sequential Experimental Design in Nano-Bio Research." Air Force Office of Scientific Research, Natural Materials, Systems and Extremophiles Program Review, Washington, D.C., January 2013.
- P.I. Frazier "Ranking and Selection with Tight Bounds on Probability of Correct Selection." INFORMS Computing Society Conference, Sante Fe, NM, January 2013.
- P.I. Frazier and J. Xie "Bayes-optimal Policies for Multiple Comparisons with a Known Standard." INFORMS Computing Society Conference, Sante Fe, NM, January 2013.
- P.I. Frazier, L. Chen and B. Jedynak, "Sequential Screening: A Bayesian Dynamic Programming Analysis." Winter Simulation Conference, Berlin, December 2012.
- P.I. Frazier, L. Chen and B. Jedynak, "Sequential Screening: A Bayesian Dynamic Programming Analysis." INFORMS Annual Meeting, Phoenix, AZ, October 2012.
- P.I. Frazier and S.C. Clark, "Parallel Global Optimization using an Improved Multipoints Expected Improvement Criterion." INFORMS Annual Meeting, Phoenix, AZ, October 2012.
- P.I. Frazier and S.C. Clark, "Parallel Global Optimization with Expensive Function Evaluations: A One-Step Bayes-Optimal Method." MOPTA, Lehigh University, Bethlehem, PA, August 2012.
- P.I. Frazier and S.C. Clark, "Parallel Global Optimization Using An Improved Multipoints Expected Improvement Criterion." Uncertainty in Computer Models 2012 Conference, Sheffield, UK, July 2012. (contributed poster)
- J. Xie, P.I. Frazier, and S.E. Chick, "Value of Information Methods for Pairwise Sampling with Correlations." Uncertainty in Computer Models 2012 Conference, Sheffield, UK, July 2012. (contributed poster)
- P.I. Frazier and S.C. Clark, "Parallel Global Optimization Using Multi-points Expected Improvement and Stochastic Approximation." 2012 CORS/MOPGP International Joint Conference, Niagra Falls, ON, June 2012.

- P.I. Frazier and J. Xie, "Bayes-optimal Policies for Multiple Comparisons with a Known Standard." INFORMS Computing Society Conference, Sante Fe, NM, January 2013.
- P.I. Frazier "Ranking and Selection with Tight Bounds on Probability of Correct Selection." INFORMS Computing Society Conference, Sante Fe, NM, January 2013.
- P.I. Frazier, L. Chen and B. Jedynak, "Sequential Screening: A Bayesian Dynamic Programming Analysis." INFORMS Annual Meeting, Phoenix, AZ, October 2012.
- P.I. Frazier and S.C. Clark, "Parallel Global Optimization using an Improved Multi-points Expected Improvement Criterion." INFORMS Annual Meeting, Phoenix, AZ, October 2012.
- P.I. Frazier and S.C. Clark, "Parallel Global Optimization with Expensive Function Evaluations: A One-Step Bayes-Optimal Method." MOPTA, Lehigh University, Bethlehem, PA, August 2012.
- P.I. Frazier and S.C. Clark, "Parallel Global Optimization Using An Improved Multi-points Expected Improvement Criterion." Uncertainty in Computer Models 2012 Conference, Sheffield, UK, July 2012. (contributed poster)
- J. Xie, P.I. Frazier, and S.E. Chick, "Value of Information Methods for Pairwise Sampling with Correlations." Uncertainty in Computer Models 2012 Conference, Sheffield, UK, July 2012. (contributed poster)
- P.I. Frazier and S.C. Clark, "Parallel Global Optimization Using Multi-points Expected Improvement and Stochastic Approximation." 2012 CORS/MOPGP International Joint Conference, Niagra Falls, ON, June 2012.
- P.I. Frazier, S.C. Clark, J. Xie, R. Waeber, and S.G. Henderson, "New One-Step Bayes-Optimal Algorithms for Global Optimization: Parallel Computing and Common Random Numbers." Air Force Office of Scientific Research Grantees Meeting, Arlington, VA, April 2012.
- S.C. Clark and P.I. Frazier, "Parallel Global Optimization Using An Improved Multi-points Expected Improvement Criterion." INFORMS Optimization Society Conference, Miami, February 2012.
- P.I. Frazier and A.M. Kazachkov, "Guessing Preferences: A New Approach to Multi-Attribute Ranking and Selection." Winter Simulation Conference, Phoenix, December 2011.
- P.I. Frazier, J. Xie (presenter), and S.E. Chick, "Bayesian Optimization via Simulation with Correlated Sampling and Correlated Prior Beliefs." Winter Simulation Conference, Phoenix, December 2011. (invited talk)
- R. Waeber (presenter), P.I. Frazier, and S.G. Henderson, "A Bayesian Approach to Stochastic Root-Finding." Winter Simulation Conference, Phoenix, December 2011. (invited talk)

- P.I. Frazier, "Tutorial: Bayesian Methods for Global and Simulation Optimization." INFORMS Annual Meeting, Charlotte, November 2011 (invited INFORMS tutorial)

- P.I. Frazier, "Indifference-Zone Ranking and Selection with 10,000 or More Alternatives." INFORMS Annual Meeting, Charlotte, November 2011

- J. Xie (presenter), P.I. Frazier, "Sequential Bayes-optimal Policies for Multiple Comparisons with a Control." INFORMS Annual Meeting, Charlotte, November 2011 (invited talk)

- J. Xie, P.I. Frazier, "Sequential Bayes-optimal Policies for Multiple Comparisons with a Control." INFORMS Annual Meeting, Charlotte, November 2011

- S.G. Henderson (presenter), P.I. Frazier, R. Waeber "A Bayesian Approach to Stochastic Root Finding" INFORMS Annual Meeting, Charlotte, November 2011 (invited talk)

- J. Xie (presenter), S.E. Chick, P.I. Frazier "Bayesian Optimization via Simulation with Correlated Sampling and Correlated Prior Beliefs" INFORMS Annual Meeting, Charlotte, November 2011 (invited talk)

- P.I. Frazier, "Indifference-Zone Ranking and Selection with 10,000 or More Alternatives." INFORMS Simulation Society Workshop, Montreal, July 2011. (contributed poster)

- P.I. Frazier, "Sequential Ranking and Selection: Tight Bounds and Large-Scale Problems." INFORMS Applied Probability Society Conference, Stockholm, July 2011 (invited talk)

- P.I. Frazier, Z. Owen, R.C. Bicalho, T.M.A. Santos, A.G.V. Teixera "Optimal Sequential Experimental Design for Stochastic Root-finding in Drug Development." INFORMS Healthcare Conference, Montreal, June 2011

- P.I. Frazier "Bayes-Optimal Methods for Simulation Optimization." Air Force Office of Scientific Research Grantees Meeting, Arlington, VA, April 2011